Next Generation Relay Services at Mars via an International Relay Network

Roy E. Gladden,[1] Greg J. Kazz,[2] Scott C. Burleigh,[3] Daniel Wenkert,[4] Charles D. Edwards, Jr.[5]
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109, USA

Nearly all data acquired by vehicles on the surface of Mars are returned to Earth via Mars orbiters – more than 1.7 terabits so far. Successful communication between the various spacecraft is achieved via the careful implementation of internationally recognized CCSDS telecommunications protocols and the use of planning and coordination services provided by NASA's Mars Program Office and the Multimission Ground Systems and Services (MGSS) Program at the Jet Propulsion Laboratory in Pasadena, CA. This modern Mars relay network has evolved since its inception in 2004 with the addition and loss of several missions, but it has fundamentally remained unchanged. Ground interfaces between the various spacecrafts' mission operations centers on Earth remain largely unique for each participant, each mission maintains its own interfaces with deep space communications networks (e.g. DSN, ESTRACK), which are similar but still unique; and relay sessions at Mars require careful ground planning, coordination, and implementation. This paper will discuss the existing architecture and consider how several technologies may be applied to the next generation of relay services at Mars. Ultimately, these are expected to lead to the implementation of a delay and disruption tolerant network at Mars, a precursor to becoming a major element in an emerging Solar System Internet.

I. Nomenclature

In this document, the following terms are used:

- The term "data" refers to binary information that is needful to be received by or sent from an entity. This term intentionally makes no assumption about the content, formatting, size, or construct of the data.
- The terms "data file" and "data product" imply that data may be assembled into a distinct construct that may be received by or sent from an entity.
- The term "relay services" implies the transfer of data from one entity to another by an intermediary who provides the services.
- The term "relay service user spacecraft" refers to a spacecraft (as used herein), that requires relay services of another spacecraft.
- The term "relay service provider spacecraft" refers to a spacecraft (as used herein), that provides relay services to another spacecraft. It is notable that a spacecraft may be both a relay service user spacecraft and a relay service provider spacecraft.
- The term "node" refers to an entity that acts as a relay service provider or a relay service user and may refer to either a spacecraft, a ground station, or other ground-based network components.
- The term "relay session" refers to a communications session between a relay service user spacecraft and a relay service provider spacecraft during which data may be transferred.
- The term "lander" is deliberately used in this document to refer to any spacecraft on or near the surface of Mars, distinct from those that may be orbiting Mars. Generally, a lander is a relay service user spacecraft but this should not be considered a strict equivalency. When specifically discussing the existing Curiosity and Opportunity rovers on the surface of Mars, the term "rover" may be used interchangeably.

---

[1] Manager, Mars Relay Network Office, and AIAA Senior Member.
[2] Technical Group Supervisor and Principal Engineer, Project Software and Information Systems Engineering Group.
[3] Principal Engineer, Project Software and Information Systems Engineering Group, and AIAA Member.
[4] Technical Group Supervisor, Instrument Operations Engineering Group, and AIAA Member.
[5] Manager, Mars Program Formulation Office, and AIAA Senior Member.

- The term “orbiter” is deliberately used in this document to refer to any spacecraft that is orbiting Mars, distinct from those that may be on or near the surface of Mars. Generally, an orbiter is a relay service provider spacecraft but this should not be considered a strict equivalency.
- The term “return-link” refers to the transfer of data from a relay service user spacecraft to the operators of the relay service user spacecraft via a relay service provider spacecraft and other intermediate network nodes.
- The term “forward-link” refers to the transfer of data from the operators of a relay service user spacecraft to the relay service spacecraft, opposite the “return-link” data flow, through a relay service provider spacecraft and other intermediate network nodes.
- Data transferred to a relay service user spacecraft is referred to as “forward-link data”. This data may be a data file.
- Data received from a relay service user spacecraft is referred to as “return-link data”. This data may be a data file.
- The term “sol” refers herein to a Martian day.

In this document, the following acronyms are used:

- AIAA = American Institute of Aeronautics and Astronautics
- APID = application identifier
- BCB = Block Confidentiality Block
- BIB = Block Integrity Block
- BIBE = bundle-in-bundle encapsulation
- BP = Bundle Protocol
- BSP = Bundle Security Protocol
- CCSDS = Consultative Committee for Space Data Systems
- CFDP = CCSDS File Delivery Protocol
- DSN = Deep Space Network
- DTN = Delay Tolerant Network
- EMI = electromagnetic interference
- ESA = European Space Agency
- ESTRACK = ESA’s deep space Tracking stations
- FHLH = first-hop/last-hop
- FTP = File Transfer Protocol
- InSight = Interior Exploration using Seismic Investigations, Geodesy and Heat Transport (lander)
- IP = Internet Protocol
- LTP = Licklider Transmission Protocol
- MAC = Media Access Control (address)
- MAVEN = Mars Atmosphere and Volatile Evolution (orbiter)
- MEX = Mars Express (orbiter)
- MGSS = Multimission Ground Systems and Services
- MRO = Mars Reconnaissance Orbiter
- MTU = maximum transmission unit
- NASA = National Aeronautics and Space Administration
- PDU = protocol data unit
- Prox-1 = Proximity-1 (protocol)
- PUS = Packet Utilization Standard
- SMTP = Simple Mail Transfer Protocol
- SSI = Solar System Internet
- TCP = Transfer Control Protocol
- TGO = Trace Gas Orbiter
- USLP = Unified Space Link Protocol

II. Introduction

NEARLY all data acquired by vehicles on the surface of Mars are returned to Earth via Mars orbiters – more than 1.7 terabits. Today, the Curiosity and Opportunity rovers from the National Aeronautics and Space

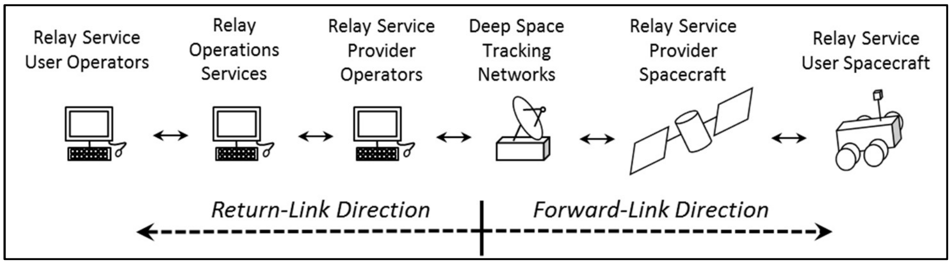

Fig. 1 General Relay Network Topology.

Administration (NASA) may utilize up to five Mars orbiters, including two from the European Space Agency (ESA), to relay this data back to Earth, which capability represents a highly successful international collaboration. Communication between the various spacecraft is achieved via the careful implementation of Consultative Committee for Space Data Systems (CCSDS) telecommunications protocols and the use of planning and coordination services provided by NASA's Mars Program Office and the Multimission Ground Systems and Services Program at the Jet Propulsion Laboratory in Pasadena, CA.

The topology of the relay network includes the proximity links between the spacecraft at Mars, the deep space links between the orbiters at Mars and their Earth-based ground stations, and the ground system links between the ground stations and the mission operations centers for the orbiters and the rovers at Mars. This topology also rightly includes the infrastructure required for the planning, coordination, and implementation of activities at each node in the network, inclusive of ground systems, deep space antennas, and the spacecraft themselves. A generalized view of this topology is illustrated in Fig. 1.

This modern Mars relay network has evolved since its inception in 2004 with the addition and loss of several missions at Mars, but it has fundamentally remained unchanged. Ground interfaces between the mission operations centers of the various spacecraft on Earth remain largely unique for each participant, each mission maintains their own interfaces with the deep space networks, which manifest as similar but still unique interfaces; and the relay sessions at Mars require careful ground planning, coordination, and implementation.

For the small, existing relay network, these one-to-one interfaces are manageable and the hands-on approach to relay planning serves the participants well. However, as the network is expected to grow and evolve in the next decade to include more spacecraft from a wider cast of participating organizations – and potentially even including human exploration components – it is useful to consider what other technologies might be needed.

This paper will discuss several of these pending technologies, which are predicted to be necessary for the next generation of relay activities at Mars.

III. Precursor Technologies

All of the following capabilities are necessary precursors to the implementation of a delay and disruption tolerant network, which further requires automated, *in situ* communications scheduling. The implementation of these capabilities is expected to be guided or informed by existing and emerging CCSDS protocol and mission operations standards, which will be mentioned in this paper.

A. Addressable Data Transfers

In the terrestrial Internet, most data communicated between machines includes an address that defines its intended endpoint. For example, when a user sends an email, the header of the email contains a construct that indicates the end recipient of the message. This email address is used by the transmission protocol to route the data through the network to the recipient. The address is meta-data associated with the content of the message, and is not typically considered part of the message itself. The transmission protocol that manages the message (either to receive it or to transfer it elsewhere in the network) reads and interprets this meta-data to recognize the recipient. In the terrestrial Internet, most machines that handle emails typically do not directly know all of the intended recipients network-wide, but instead they know enough to route the email to another machine on the network that is assumed to be closer to the end recipient. In this way, data is forwarded from machine to machine until the message is delivered.

By contrast, all of the orbiters that are currently operating as part of the Mars relay network can receive data from a lander on the surface of Mars, but the orbiters only know one destination to which to transfer the data: Earth. The orbiters receive the data as binary data and package them as if they were orbiter data.

For example, NASA's Mars 2001 Odyssey orbiter, the oldest spacecraft in the extant Mars relay network, is told by ground operators the identity of the asset on the surface of Mars with which it communicates. It uses this information to attribute an application identifier, or APID, to the data prior to transmitting the data to Earth. There is no other distinguishing marker applied by the orbiter to the data at transmission time, and it is left to ground operators to sort out the actual end recipient of the data, a matter that is complicated if the same APID is applied to data received from more than one lander. ESA's Mars Express (MEX) Orbiter handles this return-link data in a similar manner.

The newest orbiter in the network, the ExoMars Trace Gas Orbiter (TGO), operated by ESA, receives data from a surface asset during a communications session and packages all the received data as a single CCSDS File Deliver Protocol (CFDP, see Ref. [1]) data product (i.e. a file). As with the Odyssey orbiter, TGO must also be told the identity of the asset from which data will be received prior to the relay session; this information is applied to the CFDP file header. Despite having this destination information available to it, TGO does not interpret this destination data in any way, but again packages the data as if it were TGO data and transmits it to Earth. Ground processes must assemble the CFDP file on the ground, identify the intended recipient, and then route the data appropriately. NASA's Mars Reconnaissance Orbiter (MRO) and the Mars Atmosphere and Volatile Evolution (MAVEN) orbiter handle return-link data in a similar manner.

In all cases, ground processes take an active role in the transfer of data to the end recipients, typically the mission operators or scientists operating the relay service user spacecraft. These ground processes are responsible for de-packaging the data as orbiter data to access the lander data within it, handling any CFDP processing of the data, typically to remove the CFDP packaging; and then identifying the end recipient and delivering it. The situation is complicated by the use of multiple ground station networks, such as NASA's Deep Space Network (DSN) and ESA's European Space Tracking (ESTRACK) network, which requires a multitude of implementations to support the various paths through which the data may flow.

Despite the apparent simplicity of this approach, which serves the current needs of the network quite well, there are several fundamental weaknesses:

1. The orbiters must be informed, prior to the relay session, of the identity of the relay service user with whom a relay session will occur. This has led to an architecture where the orbiter must always be configured either to initiate the relay session itself or to be proactively commanded to expect to receive a signal from a relay service user spacecraft at specific times. (See also Section D for more information on this operational paradigm.)
2. The orbiters cannot support receiving data from a local relay service user spacecraft and then forward it to another relay service user spacecraft at Mars. This prohibits supporting orbiter-to-orbiter, lander-to-orbiter, orbiter-to-lander, and lander-to-lander relay services (i.e. without Earth as an endpoint).
3. Ground processes that handle the data, as received from the five orbiters, are unique for each orbiter. This has required the implementation of unique ground processes for de-packaging the relay data from each orbiter's data, identifying the intended recipient of the data, and then transferring that data.

To achieve addressable data transfers at Mars, the use of the Delay-Tolerant Networking (DTN) architecture's Bundle Protocol (BP), a CCSDS standard, is proposed (Ref. [2]). Future orbiter missions could then, at the time of receipt of data from a relay service user spacecraft, identify the destination and automatically determine the next step to delivering the data to its intended recipient. Similarly, an orbiter's own science instruments could identify the recipient as the science teams who operate the instruments. In today's network, the next node to which the orbiter should transmit these packaged data would still be to "Earth", but then ground stations could receive the data and directly transfer it to the recipient without the need to de-package and otherwise process the data. In addition, data transfers to destinations other than Earth could be supported.

Operational experience with the DTN architecture suggests that it will be applicable to a Mars relay network; DTN was demonstrated on the EPOXI spacecraft in deep space in 2008 (Ref. [26]) and has been operational on the International Space Station since 2016 (Ref. [27]). However, one challenge in implementing this capability is that the entire network, or at least several connected parts within it, would need to understand the addressing structure within BP. In the current space community, many spacecraft are built using propriety technologies or formatting. This is a challenge that directly inhibits adoption. To help overcome this, a DTN "first-hop/last-hop" (FHLH) mechanism is being developed which will enable a BP-cognizant spacecraft to communicate with a non-BP-cognizant spacecraft using legacy links.

B. Custody Data Transfers

In the terrestrial Internet, many protocols exist to manage the transmission of data across a network. These protocols often include a method to ensure that the data is completely transferred without error. Such protocols include the common File Transfer Protocol (FTP) (Ref. [3]) and the Simple Mail Transfer Protocol (SMTP) (Ref. [4]). In principle, reliable data transfers are effected when the receiving entity informs the sending entity that it has received the transmitted data. This can be verified by the use of checksums and other error detection schemes. When a transfer has concluded successfully, the sending entity can "forget" about the data, and the receiving entity then takes responsibility for delivering the data to the end recipient. This is called a "custody transfer".

In the current Mars relay network, there exists no notion of a custody transfer. When data that has been generated by a lander is transmitted to an orbiter, the data is received as simple binary data. No specific format or content of the received data is assumed, as mentioned in Section A. Usually, the data is reliably transferred by using the CCSDS Proximity-1 (Prox-1) Protocol (Ref. [5-7]),[6] which ensures that the data is moved reliably as a bitstream.[7] In this way, no data is lost or duplicated in the transfer. However, the orbiter doesn't take strict custody of the data because it does not understand its content. For example, a lander may send several images during a single transfer, and the orbiter has no mechanism to detect the boundaries of the image files within the data it has received.

The lack of knowledge regarding the structure of the received data prohibits the use of a true custody transfer from the lander to the orbiter. The orbiters simply turn the data around and transmit the data to Earth without regard for its internal data boundaries. Reliability on this direct-to-Earth transfer is ensured on a few of the existing spacecraft using different retransmission schemes, but even those implementations are not common across the orbiters. TGO, for example, uses the CCSDS Packet Utilization Standard (PUS) Service 13 (Ref [9]) to ensure reliable transfers from the orbiter to Earth, but this again is applied at the level of the orbiter data product, which can include many data products as received from a lander. MRO uses a proprietary retransmission scheme that suffers from the same limitations.

These retransmission schemes, even when applied, are not 100% effective, and occasionally there remain gaps in the received data. Data may also be lost in transmission from the ground stations to the mission users, but this is extremely rare.

In the end, incomplete data may be received by the lander operators. These operators may then choose to command the retransmission of the missing data from the lander on the surface of Mars, thereby starting the whole chain of communications over again for that data, or they may choose to accept the loss. Due to the inherent delays in the network, sometimes the original data is no longer available onboard the lander and the data is unrecoverable.

If a true custody transfer of the data were to be implemented between the various nodes in the network, then the data could be moved off the lander on the first reliable transmission and the lander could then "forget" about it immediately. The orbiter would then take full responsibility for transmitting the data to Earth and could do so reliably. Upon receipt, the ground stations would then take full responsibility for transmitting the data to the next recipient, etc.

This notion of a custody transfer requires that the data be packaged as recognizable Protocol Data Units (PDUs), and that each node in the network recognize the boundaries of those units. It also relies upon the existence of a reliable transmission of the data between every node in the network. Here again, the BP could be applied to answer the question of packaging. The reliability question could be managed on the direct-to-Earth transfer using a variety of methodologies, but it is proposed that the DTN Licklider Transmission Protocol (LTP) (Ref. [10]) be applied for reasons that will be further explained in Section F.

Several problems are identified relating to custody transfers in deep space applications. First among them is that a reliable transmission of the data may require several retransmissions in order to acquire data that was lost on the first transmission. A reliable exchange between a pair of nodes that are in close proximity to each other (where the transit times are low, such as between a lander and an orbiter) may not be greatly affected by this. However, in deep space applications, the time it takes for the radio signals to transit between the nodes (such as between an orbiter and the Earth) may cause a significant delay when attempting to assemble a complete data product.

For the operators of both the Curiosity and Opportunity rovers on the surface of Mars, this delay may be unacceptable. Both operations teams typically generate commands for the next sol's activities after receiving data from the prior sol's activities. This is often necessary because the rovers often physically move from day to day and the local environmental conditions on the surface of Mars may prove hazardous to the long-term health of the rovers, or the missions themselves may have a short design lifetime. In both cases, the mission operators are willing to accept receipt of their data even if there exists significant gaps within it. To them, some data is better than no data. For this sort of operational scenario, it may be advisable to use a protocol like CFDP that can deliver partial data products incrementally as segments are received.

In the other direction, when transmitting commands from Earth to the Curiosity and Opportunity rovers, the late receipt of command products onboard the rovers might prove risky to the mission, as commands are quite often time-dependent. In addition, the commands are typically not even available to send to the rovers until very close to their

---

[6] Note that the CCSDS Proximity-1 Protocol is expected to be replaced by the Universal Space Link Protocol (USLP), as in Ref. [8].

[7] Note that the current Mars Relay Network implements the Proximity-1 protocol as a reliable bitstream. However Prox-1 does have a provision for reliable packet transfer, as well, which enables the accountable transfer of data units across the network.

execution time, with little occasion for retransmission if the commands are lost anywhere in the transfer path. Thus, in the return-link direction, the rover operators will accept partial data returns, but are not content with delays in the return; and in the forward-link direction, the operators will accept only an on-time and complete transfer.

The classical method for assuring reliable data delivery in a network is retransmission, as is performed by LTP over links characterized by extremely long signal propagation delays. When the resulting final delivery latency is unacceptably high, "erasure coding" can be used instead, whereby a data item is segmented into some number of data blocks and a number of "parity blocks" are computed from the content of the data blocks such that reception of a combination of data blocks and parity blocks can suffice to recreate the original data item (notably without retransmission) even if multiple data blocks are lost in transit. The number of parity blocks computed from the data blocks can be arbitrarily increased until the probability of data item recreation failure drops to nearly zero.

C. Multi-Channel Support and Frequency Agility

Terrestrial computer-to-computer interfaces typically function using an exponential backoff approach to ensure that a message is passed successfully between computers. The target computer is identified by its Media Access Control (MAC) address, which is attached to the message. As many computers may exist on the same transmission medium (either a wired network or an over-the-air wireless network), a computer will wait until the medium is clear before attempting to send a message. Computers on the medium that are not the intended recipient ignore the message. If such a "data collision" occurs at one attempt, the computer will wait to try again. The duration it waits exponentially increases until the medium becomes clear for the data to be transferred.

One way to minimize the probability of a data collision on the network is to reduce the number of computers sharing the same transmission medium. In a wired network, this is often done by using a network switch, which effectively isolates several machines on a network into smaller sub-networks, which passes data only when data is detected on one side of the switch that is addressed to the other side of the switch.

Another way to ease congestion on a network is to chop large messages into smaller ones using a technique called segmentation. These smaller segments can be routed through a variety of pathways until they reach the end destination where it can be reassembled into the final message. This is a primary function of the TCP/IP protocol (Ref. [11, 12]), and facilitates the insertion of small messages into the network from a variety of sources in the presence of very large messages, all of which may have different destinations.

Some wireless devices can take advantage of different frequencies on which they can transmit data. For example, some wireless routers for use in the home can support transmissions on multiple frequencies, the most common being 2.4 GHz and 5 GHz. Wireless devices that use such a router can detect if one particular frequency has a lot of interference on it, presumably from network traffic from other devices, and then opt to use the other frequency in the hope that there will be fewer data collisions during data transfers.

In the Mars relay network to date, data collisions have largely been avoided by the sparsity of nodes in the network. The Curiosity and Opportunity rovers, for example, are on opposite sides of Mars, and thus can never communicate with the same orbiter simultaneously. However, in late 2018, the Interior Exploration using Seismic Investigations, Geodesy and Heat Transport (InSight) lander will land close enough to the Curiosity rover that they will share visibility to the relay orbiters over 87% of the time (the remaining 13% of non-shared view periods are low-elevation, horizon-skirting view periods that are not generally suitable for performing data transfers).

The Proximity-1 Protocol designates that relay sessions are to be established between two spacecraft on what is called "Channel 0", or the "hailing" channel. Data is addressed internally with a spacecraft identifier, which is functionally equivalent to a MAC address. Just as in the terrestrial Internet, if a message is received by a spacecraft's radio that contains a spacecraft identifier that does not match the one for that spacecraft, it ignores the message.

After a relay session is established on Channel 0, the Curiosity rover can accept a directive from the orbiter to move to a different channel. When communicating with MRO, the Curiosity rover typically uses what is called "Channel 2." This is the "working" channel between MRO and Curiosity, which is preferred over Channel 0 due to electromagnetic interference (EMI) onboard MRO by several of its science instruments. Communications between the Curiosity rover and all other orbiters in the network remain on Channel 0 after the hail is established because they do not suffer from this EMI.

By contrast, the Opportunity rover does not have the ability to switch working channels and thus always communicates on Channel 0. The Odyssey orbiter and the InSight lander both carry the same model of radio as the Opportunity rover and therefore bear the same complete reliance on Channel 0. All three future missions planned for launch in 2020 will have a similar mix of capabilities.

This dependency on Channel 0 for establishing a relay session and for maintaining the relay session (in many cases) is a double-edged sword. On the one hand it ensures interoperability between the assets in the near-Mars environment, but on the other hand it means that only one communications session may be active at any given time in a single view. Even in the MRO-Curiosity case where the working channel is different than the hailing channel, it is still desired to keep the hailing channel clear in case the link is interrupted and the session has to be re-established.

As it is, the operators of the Curiosity rover and the InSight lander have developed strategies for sharing the relay resources that are available to them. This is done by coordinating *a priori* which lander gets to talk to which orbiter as a function of mission phase and time of day. This will ensure that only one relay session occurs at a time, but at the potential cost of reducing the overall throughput for both missions.

As implemented, the current relay sessions are largely “all or nothing” affairs, where they are pre-scheduled to occur on the orbiters and the landers; the orbiters take no notice of the structure of the received data, thus prohibiting multiplexing of received data from multiple sources. Careful effort is expended during planning to ensure that only one relay session occurs at any given time within a given view, which altogether avoids the possibility of data collisions that are so common in the terrestrial Internet. However, this approach is only practical due to the small size of the network and won’t remain practical as the network increases in size.

Given the hardware on the existing assets at Mars, these problems can’t be overcome in the near-term. However, future missions could ensure that they have the ability to be more frequency agile, or even carry the ability to attempt communicating on multiple frequencies at the same time. In principle, spacecraft at Mars could employ the same techniques used by terrestrial wireless devices to detect a busy frequency and switch to another one. Fragmenting large data transfers into appropriately-sized PDUs could further increase the flexibility of the network, reduce the likelihood of data collisions, and maximize the transport efficiency of the network. This has implications for the design of the relay asset’s avionics, which, as mentioned in Section A would need to support addressable data transfers.

D. Demand Access and Always-On

Terrestrial computers are designed so that a computer is effectively always able to send data if it’s plugged into a network (notwithstanding the potential for data collisions as described in Section C) with an assurance that receiving computers on the network are listening and able to accept the data. The cost of maintaining a live and active internet connection is small relative to the cost of powering the computer, so typically the network connection is maintained as a “live” connection.

In wireless applications, this is not necessarily the case. Wireless devices, including the ubiquitous cellphone, may only power on their transmitters when establishing a connection or when data needs to be sent. After connectivity is established with the nearest network node and any data that needs to be sent has been transferred, the devices power down their transmitters, and only occasionally ping the network to ensure connectivity remains available. If connectivity is lost, such as if the device has moved out of range, the device may seek for another network connection, continually searching until a connection can be established.[8] In this manner, the transmitters remain active (and drawing power) only when sending data or otherwise ensuring connectivity.

When data do need to be sent, wireless devices first need to establish connectivity in a “session” with the receiving radio. In the case of a wireless device in a home, as mentioned in Section C, this session is typically established with a wireless router, which itself is plugged into a wired network and continuously powered. In this example, the router is considered to be “always on”, constantly receptive to a signal from a wireless device with the ability to service multiple devices in a time-shared manner.

In the current Mars relay network, the Opportunity and Curiosity rovers do not enjoy the benefit of an “always-on” network. Even though an orbiter may be physically visible, they are not receptive to a signal from the lander unless they have been actively scheduled to be so. In addition, the current network is designed so that reliable, full duplex communications sessions are only initiated by the orbiter. Thus, a lander with data to be sent can only send that data to the orbiter if the orbiter first establishes the relay session.

This approach requires meticulous planning by the spacecraft operators. The orbits of the orbiting spacecraft need to be carefully predicted and/or controlled. This is required so that view periods between the orbiters and the landers can be forecast with sufficient accuracy and lead-time to facilitate the commanding processes required to schedule the relay sessions on both ends. For low-altitude orbiters, this lead-time is measured in weeks, which constrains the typical last-minute planning approaches utilized by the rover operators, who respond to realtime conditions at the landing site and prefer to do daily commanding of the vehicles.

The amount of data that can be communicated during each relay session is predicted by the rover operators. The amount of science data acquired by the vehicles during their missions is limited by this estimated available data return. To complicate matters further, the exact performance of each relay session when compared to the predicted performance regularly varies by as much as 20% or more. This implies that the data management scheme onboard

---

[8] This searching for connectivity explains why a cellphone’s battery drains so quickly when it is out of range of the network.

the rovers needs to be flexible enough to accommodate both an over-performance of a relay session and an under-performance.

If one were to shift the operational paradigm so that the orbiters at Mars acted analogously to their wireless router counterparts on Earth, several changes would need to be made. The orbiters would need to have their onboard radios constantly powered so that they could receive a signal at any time. For some of the existing orbiters, the radios are only powered on at the time of a relay session. Also, the orbiters would then need to be reprogrammed so that a relay session could be initiated by a landed asset, and these at any time (i.e. "demand access").

Presently, the high reliance on navigation predictions for the orbiters when planning relay sessions further constrains the relay network architecture. For example, proactive scheduling of relay sessions must account for some nominal uncertainty in the orbits. The rover operators, knowing that power and time are precious commodities onboard the rovers, will typically schedule a relay session to occur when the orbiter is predicted to be well above the horizon. This conservatism ensures that navigation errors do not cause the relay session to slip from view, but it has the side effect of underutilizing the potential throughput of a given view period. While data throughput at low elevations can be limited, some data can usually still be transmitted.

Implementing a demand access approach with the orbiters would help mitigate the problem of the orbiter navigation uncertainty. The rover could ping the orbiter for access at some reasonably close time to the expected view period, and, if connectivity is available, the rover could begin the relay session immediately. If no response is received from the orbiter, the rover could wait and try again later. This strategy would both reduce the dependence on the accuracy of the predicted orbit, as well as improve the overall data throughput. (Note that the use of the DTN LTP protocol can further improve orbital contact utilization; transmission can begin at the earliest possible moment and end at the last, because any data that are lost due to low signal quality are automatically retransmitted as soon as possible.)

Of course, the issue of needing onboard data management flexibility doesn't go away in a demand access environment. However, such a network would allow the rover to autonomously initiate an otherwise unscheduled relay session with an orbiter to return data to Earth that may not have been transmitted when it was originally expected to be transmitted. With low-altitude orbiters, communication opportunities are sparse and short, so some wisdom would need to be applied to the rovers to not waste time and power attempting to communicate with an orbiter that won't be there for long periods. However, if higher-altitude orbiters were to be sent to Mars, continuity of coverage could become a reality and a demand access approach could become more meaningful to relay users.

This concept, coupled with that of frequency agility as discussed in Section C, would be a game-changer for the types of missions that could be supported at Mars. Many users could use the same relay assets simultaneously, and other mission types previously unsupported could be added, such as constellations of very small satellites that have little to no means to perform navigation and/or to communicate directly to Earth.

E. Node Richness and Trunk-Lines

In the terrestrial Internet, a multiplicity of connections between many computers provides many routes for data to be moved between widely separated machines. Even with wireless applications, usually only one or two hops are necessary before data is inserted into the wired Internet where this multiplicity of pathways facilitates the rapid transfer of data to its recipient. The network is node rich and the infrastructure is generally oversized for the amount of data transferred.

Consider the network topology shown in Fig. 2. The diagram includes a representation of the types of networked devices in the home of the lead author of this paper, in a general network diagram. Note that nearly all of the devices are wireless, with most connecting to the range extender noted in the lower middle portion of the diagram. It is notable that despite the

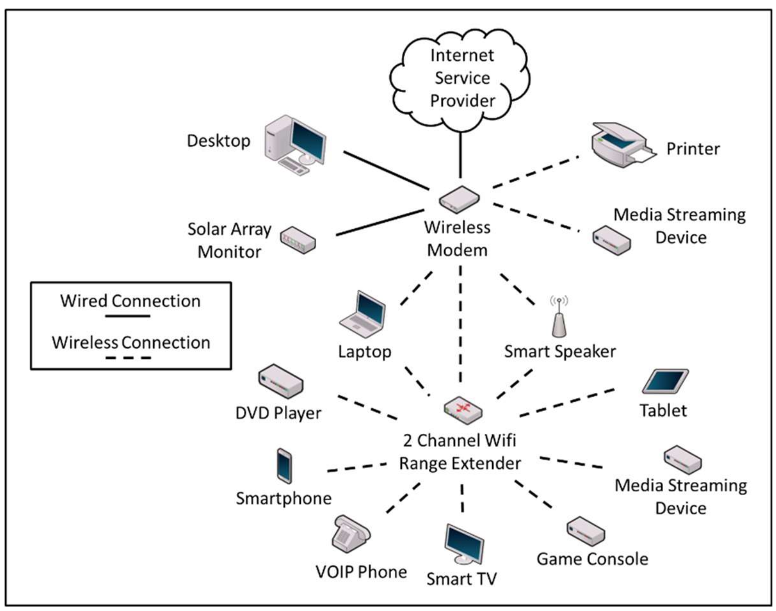

Fig. 2: Example Home Network Topology.

plethora of devices and device types, all of them ultimately gain access to the Internet via the single wireless modem, which acts as a trunk line.

Each device independently operates on the network and negotiates with the modem and range extender. Data collisions, as described in Section C, undoubtedly occur quite frequently, especially when media is being streamed, and yet the internet functions seamlessly with no directed coordination among the devices and without thought on the part of the users (except when streaming movies stall!).

The current Mars relay network is similarly "node poor." With five orbiters at Mars (only three of which are presently, actively, and regularly used to provide relay services), there aren't many paths that data from the existing rovers can take to be returned to Earth. Each orbiter acts as an independent trunk line for the rovers, but, as mentioned in Section D, relay sessions with the orbiters have to be deliberately scheduled.

Fortuitously, the two rovers are on opposite sides of Mars and thus have not had to coordinate relay opportunities between them. However, a larger community of relay service users is expected to arrive at Mars in the next few years. The existing relay network will then be stretched to provide relay services to these additional assets. If even more assets were to be sent, such that there were more than a dozen relay users, the existing network would simply be unable to meet all of their needs.

In contrast, consider if a network of three telecommunications satellites sharing an equatorial orbit at 6,000 km altitude were to be sent to Mars. Together, they could provide continuous relay coverage to landing sites up to 60 degrees latitude (north and south). Such satellites that support multi-frequency communications (as in Section C) and demand access features (as in Section D) would be able to support a variety of missions on the surface as well as lower-altitude orbiters. They could serve as trunk lines, providing relay services to dozens of assets with a robust line of communication back to Earth. In this way, the relay network at Mars could look very much more like that in Fig. 2, with a variety of users gaining access to the network that isn't much broader in scope than that in a typical home.

IV. Delay and Disruption Tolerant Networks

F. Data Loss and Latency

Terrestrial networks experience almost no delay when communicating from one node to another. Transmissions occur at the speed of light and all nodes, wired or wireless, are physically within much less than 1 second from each other. (Even the one-way transit time from the Earth to the Moon is less than 1.3 seconds.) As mentioned in Section C, even when many data collisions are encountered on the network, a sending machine merely waits a matter of seconds to try sending data again.

In deep space applications, the two-way light time between two nodes in the network may be measured in minutes or hours. If the same principles were applied to deep-space communications as are applied in the terrestrial network, a data collision wouldn't be encountered and detected by the transmitting node until a two-way light time duration later. This is further limited by the geometry of the environment (eg. orbiters can be occulted by planets, landers can rotate out of view, etc.).

Due to these limitations, the contact times between Earth-based transceivers are proactively and carefully scheduled with the orbiters in the Mars relay network. Relay sessions between the orbiters and the landers at Mars may utilize reliable data exchanges (such as the Proximity-1 Protocol mentioned in Section B) to ensure that data is completely transmitted from the lander to the orbiter. However, the NASA orbiters then turn that data around and transmit the data to Earth without assurance that they will be completely received on Earth successfully. When the data is received on the ground, some of the operations teams for the orbiters use ground-based algorithms to detect transfer frames that were lost in transit, and can queue up spacecraft commands to cause the orbiter to retransmit the missing frames. Today, this frame detection is performed on the orbiter frames, which may contain any number of (or only a portion of) data products as received from the lander.

The ExoMars TGO does utilize the Packet Utilization Standard (PUS) Service 13 to ensure that data products from the orbiter are received completely on Earth. This is an automated process where missing frames are detected in realtime and spacecraft commands are sent back to the orbiters to retransmit them. Both approaches incur additional latency in the return of the data to the lander operators when frames are missing, and both have processes in place to deliver partial data sets when a complete data set is not available after the first attempt to transmit it.

This quick data delivery, even of partial data sets, is highly desired by the lander operators, as they attempt to construct the command load for the next sol's operations using whatever data they can get that informs them of the state of the rover on the surface of Mars. These operators have learned to accept the partial data sets and then to expect a more complete delivery of the data hours or days later.

Note that in all cases, none of the orbiters manage the data as received from the landers as if they were data products with known boundaries, as described in Section B. The data is received from the landers, and the orbiters manage them as a series of bits, without regard for packet, frame, product, or other data constructs within the data itself. Thus, at the end of a single relay session, the complement of data to be returned to Earth may include partial data sets from the originating source. Interestingly, this is consistent with how data is transmitted across the terrestrial network using TCP/IP, where messages are fragmented into smaller components and then transmitted piecemeal across the network, to be reassembled by the receiving node.

In the case of the Mars relay network, accountability for the complete receipt of data products as produced by the rovers on the surface of Mars is performed by the lander operators. Typically, data are held onboard the rovers until they have been confirmed as received, and then explicitly "released" via commands sent to the rover. The notion of a custody transfer, as described in Section B, would cause this accountability to occur across each leg of the network. If this were to occur, then the lander, once it had confirmed successful transfer of a data product to an orbiter, could rely upon the orbiter to complete the transfer and immediately release the data from its memory space, freeing that resource for follow-on operations.

The use of the Bundle Protocol would allow a sending entity to fragment its own data into small transfer units for handling through the entire network. In this way, though each bundle may consist of parts of one or more original data products, at each node in the network the data is transmitted reliably as a distinct and known unit, without the need to assemble or further fragment it at any node into some other construct for ease of transmission. The bundle becomes the atomic unit of the transfer across the entire network, much as the packet is the atomic unit of the transfer in IP. When using Ethernet, the maximum transmission unit (MTU) size of a frame is 1500 bytes, but in the Bundle Protocol, the MTU is limited only by local operational constraints, if at all. Bundles that are hundreds of megabytes (even gigabytes) in size may be transmitted routinely under favorable conditions.

In practice, however, transmission contact opportunities are of limited length and bundles normally should be fragmented in such a way that the last byte transmitted is the last byte of a (possibly fragmentary, but still routable) bundle. Further segmentation is typically performed at the "convergence" protocol layer that underlies the bundle transmission layer, by protocols like LTP that are aware of the limits on frame size at the underlying space link layer. Because different convergence-layer mechanisms may be operating on different "legs" of the end-to-end data path, this convergence-layer segmentation and the necessary subsequent reassembly is performed separately – and possibly in quite different ways – at multiple forwarding points on the path. Reassembly of the entire original bundle occurs only at the final destination since each fragmentary bundle can theoretically take a different route through the network.

A network could be configured to utilize a delay- and disruption-tolerant protocol for communicating data across the network. Such a protocol would address the issue of interruptions in connectivity. In the Mars relay network, these interruptions many take many forms, such as: the orbiter is not visible from the lander, the lander is not able to talk to the orbiter even when visible, the orbiter is not visible to Earth, a tracking station on Earth is not configured to communicate with the orbiter, the orbiter is not able to transmit to Earth even when visible, etc. DTN empowers an enabled network to self-determine the quickest path for data to be returned to Earth. In the presence of only a few nodes, such as is present in the Mars relay network, the quickest path for lander data that can't be directly transmitted to Earth is always obvious: transfer to the orbiter, transfer to Earth, transfer to the operators. In a more connected network where more nodes exist, the pathway may not always be so obvious.

For example, consider the case where there are two orbiters at Mars, as shown in Fig. 3. In this example, assume that the lander may be able to see Orbiter 1 "now", which has the ability to transfer data to Earth at a low data rate, and Orbiter 2 may be visible at a later time and has the ability to transfer data to Earth at a high data rate. Including the ability to cross-link the data between the orbiters, there are a variety of paths that can be used to transfer the data back to Earth, only 3 of them being shown in the figure.

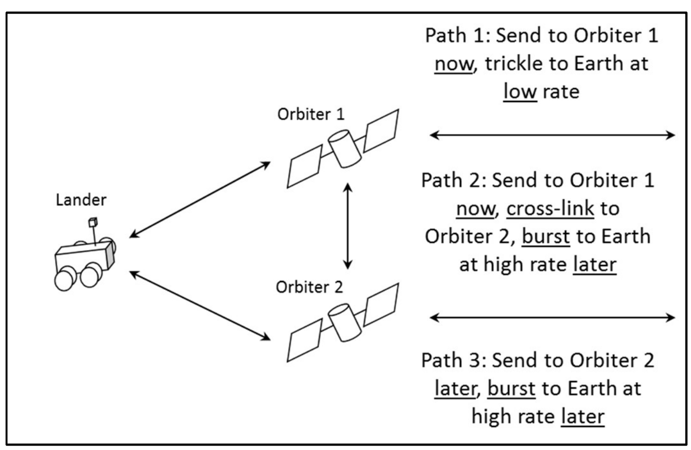

Fig. 1 Data Return Options (Contact Graph Routing)

Should the lander choose to send the data immediately to Orbiter 1, it may be that Orbiter 1 could then be confronted with the choice to send the data back to Earth directly at its low data rate, or to cross-link the data to Orbiter 2 so it can complete the transfer in its stead. It may be that Orbiter 1 can only support a low direct-to-Earth data rate, but Orbiter 2 is known to support a higher rate. But it may also be the case that Orbiter 2 doesn't have time scheduled with Earth receivers.

The use of the DTN protocol onboard the lander and the orbiters would enable each node in this network to calculate the quickest path for the data. Thus, the lander may choose to transfer the data to Orbiter 1, assuming that the quickest path for the data to be transferred is at Orbiter 1's low rate (Path 1). Orbiter 1, however, may then choose to cross-link the data to Orbiter 2 so that it may transmit the data at a higher rate at a later time (Path 2). Alternatively, the lander could choose to wait until Orbiter 2 is visible and send the data then (Path 3). The calculation to determine the quickest path should be relatively simple, but it requires a knowledge of the state and availability of every node in the network. In a DTN, this knowledge is called a "contact graph" and the act of determining the fastest route for data is called "contact graph routing." The proper application of routing in every node in the network should theoretically always minimize the data return latency.

G. Predictability and Prioritization

In principle, a DTN-enabled network could service many users in the presence of long transmission times. However, several issues manifest in the use of DTN, namely the predictability of time of data receipt and the need to perform partial data deliveries.

The current rover operators are accustomed to knowing when their data will be returned to Earth. The time it takes data to traverse the path in the current Mars relay network from the rover to the orbiter to the ground stations and then to the workstations of the rover operators is reasonably predictable. In a DTN-enabled network where there are many paths and many nodes, this predictability may be more elusive. However, software has been developed and successfully demonstrated that uses contact plan information – the same information that is used to compute routes through the network – to develop a range of potential delivery times for a bundle of a given size, source, destination, and transmission time, with the probabilities associated with each delivery time. This "bundle delivery time estimation" algorithm cannot provide a deterministically certain time of data arrival, but it can at least be used to alert network operators to potentially unacceptable latencies so that remedial action can be taken in advance.

One such remedy is simply the modification of the contact plan, scheduling additional contacts – possibly at the expense of previously scheduled contacts – or lengthening contact intervals.

Another possible remedy is providing additional forwarding advantage to the subject bundle itself, by revising the bundle's class of service. BP running over DTN supports three classes of service, namely bulk, standard, and expedited; which are commonly implemented as levels of priority in the assembly of forwarding queues. Applying these classes of service, a high-priority (i.e. expedited) message may be transferred through the network with lower latency than other lower-priority data. This would directly support the need of the lander operators to receive their critical telemetry and other data to support next-sol planning, while allowing the less critical data to be returned with greater latency.

As previously mentioned in Section B, in today's Mars relay network, data acquired during a single relay session are typically handled as a single monolothic data product when being transferred back to Earth. This prohibits the fragmentation of the data into high- vs. low-priority data. The implementation of BP on the rover asset, then, could be brought to bear on the rover side to package relay data into smaller fragments. These fragments might collectively represent one data product (such as a large picture) or might include many data products (such as critical, captured spacecraft telemetry). The use of BP allows the network to be concerned about the priority of the data, and not the content, which is controlled and abstracted by the network user.

This methodology partly addresses the desire to receive some of the data, particularly the most critical parts of it, with as low a latency as possible. While not directly addressing the goal to receive partial data products in the midst of an incomplete transfer, this methodology does effectively reduce the likelihood that this will become a practical issue.

H. Encryption

The BP does also support features that allow a user mission to encrypt their data. This is one concern for user missions when confronted with the possibility of exchanging data through a relay network where not all nodes in the network may be under the control of the organization that sponsors the mission.

The DTN architecture addresses this concern in two main ways. First, each bundle may contain specialized extension blocks that implement the Bundle Security Protocol (BSP). These blocks are of two types:

1. Block Integrity Blocks (BIBs) carry cryptographic signatures computed from the content of the various blocks of a bundle, typically the primary block and payload block. These signatures enable a receiving node to determine immediately whether or not the content of the block to which the BIB pertains (i.e. the "object" of the BIB) has been altered in any way subsequent to attachment of the BIB.

2. Block Confidentiality Blocks (BCBs) carry information that describes the manner in which the various blocks of the bundle – typically the payload – have been encrypted. This information enables the bundle's destination node to decrypt the encrypted block. Note that block encryption remains in effect not only while the bundle is in transit from one node to another but also while the bundle is "at rest" at a forwarding node while awaiting a transmission opportunity.

Taken together, the BSP blocks ensure that it is safe to forward data through a node operated by a given space agency when the source and destination of the data are some other agency altogether. A BCB applied to the payload ensures that the bundle's application information won't be disclosed to an unauthorized recipient, and a BIB applied to the payload block enables receiving nodes to detect any attempt to tamper with the metadata characterizing that information.

A second mechanism for securing data in a DTN-based network is bundle-in-bundle encapsulation (BIBE). BIBE enables an entire bundle to become the payload for a second, encapsulating bundle with a different destination. The source and destination of the encapsulating bundle may be at the entrance to and exit from a particularly hazardous region of network topology. This enables the entire encapsulated bundle to be encrypted during its transit through that dangerous space, affording a high degree of protection from traffic analysis.

I. Challenges

The challenge of implementing BP over DTN is not necessarily a technical one, but rather one of adoption. The requirements on the avionics and the telecommunications systems of the spacecraft are well-understood, having been demonstrated in flight experiments, but existing Mars missions did not have such principles in mind (and especially not as relevant requirements) when they were built. In addition, it is unfortunate that none of the orbiters in the existing Mars relay network are capable of being retrofit to support these new methodologies due to hardware restrictions, specifically in their implementations for onboard data handling.

In order for a true DTN-enabled network to be implemented at Mars, future relay service provider spacecraft would need to be designed with the principles described in this paper in mind. Methodologies for the regular distribution and refreshing of contact graph schedules onboard each of the nodes need to be developed, including via international entities. These methodologies, and others relating to network management, are currently being developed and should be available long before the next generation of DTN-capable spacecraft are launched. Once done, many network architectures could be manifest, from a distributed and many-node implementation where each node has their own connectivity to Earth, to a spoke-and-hub network where many nodes could rely upon one or a few primary nodes to transfer data back to Earth.

V. Conclusion

The technologies outlined in this paper may be implemented in any communications network that struggles with long-leg communications, whether at Mars or at other locations around the Solar System. Ultimately, it is desired to build a unified architecture that is robust to the addition and loss of nodes, both as relay service users and as relay service providers.

Beginning with just one orbiter implementing the full complement of the principles outlined in this paper, the next generation of relay capabilities at Mars can begin to be put into place. Once that occurs, a variety of mission types that heretofore have been discounted at Mars can be more achievable: small satellites, cubesats, netlanders, weather stations, mini-rovers, climbers, diggers, balloons, drones, swarms, etc. Each of these would become viable mission types, presumably alongside human explorers who would be equally served by investments in the relay infrastructure at Mars.

By successfully implementing these technologies in the next generation of spacecraft, the introduction of additional nodes would be enabled without regard to mission type or the sponsoring organization, including both government and private entities; and would provide a far more robust infrastructure than is presently available. Ultimately, the objective is to construct a true interplanetary network that functions as seamlessly as the terrestrial Internet. As more and more actors look to explore and exploit space, a unified relay network could become a critical and enabling component.

VI. Acknowledgments

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. © 2018 California Institute of Technology. Government sponsorship acknowledged. The authors would like to thank Neil Chamberlain for help in reviewing this paper.

VII. References

[1] Consultative Committee for Space Data Standards (CCSDS), "File Delivery Protocol" (CFDP), CCSDS 727.0-B-4, http://www.ccsds.org, Jan. 2007.

[2] Consultative Committee for Space Data Standards (CCSDS), "Bundle Protocol" (BP), CCSDS 734.2-B-1, http://www.ccsds.org, Sept. 2015.

[3] Internet Engineering Task Force (IETF), "File Transfer Protocol" (FTP), RFC 412, Nov. 1972.

[4] Internet Engineering Task Force (IETF), "Simple Mail Transfer Protocol" (SMTP), RFC 7504, June 2015.

[5] Consultative Committee for Space Data Standards (CCSDS), "Proximity-1 Space Link Protocol – Physical Layer," CCSDS 211.1-B-4, http://www.ccsds.org, Dec. 2013.

[6] Consultative Committee for Space Data Standards (CCSDS), "Proximity-1 Space Link Protocol – Data Link Layer," CCSDS 211.0-B-5, http://www.ccsds.org, Dec. 2013.

[7] Consultative Committee for Space Data Standards (CCSDS), "Proximity-1 Space Link Protocol – Coding and Synchronization Sublayer," CCSDS 211.2-B-2, http://www.ccsds.org, Dec. 2013.

[8] Consultative Committee for Space Data Standards (CCSDS), "Unified Space Link Protocol" (USLP), CCSDS 732.1-R-2, http://www.ccsds.org, Feb. 2017.

[9] Telemetry and Telecommand Packet Utilization Standard (PUS) Service 13, "Large Data Transfer Service," ECSS-E-70-41A, Jan. 2003..

[10] Consultative Committee for Space Data Standards (CCSDS), "Licklider Transmission Protocol" (LTP), CCSDS 734.1-B-1, May 2015.

[11] Internet Engineering Task Force (IETF), "Transfer Control Protocol" (TCP), RFC 793, Sept. 1981.

[12] Internet Engineering Task Force (IETF), "Internet Protocol" (IP), RFC 791, Sept. 1981.

[13] O. Reboud, et al, "An Interplanetary and Interagency Network - Lander Communications at Mars", AIAA International Conference on Spacecraft Operations (SpaceOps), Heidelberg, Germany, May, 2008.

[14] D. Allard. R. Gladden, "Mars Relay Operations Service (MaROS): Managing Strategic and Tactical Relay for the Evolving Mars Network," 10.1109/AERO.2012.6187115 , IEEE Aerospace Conference, Big Sky, MT, 2012.

[15] C. Edwards, B. Arnold, D. Bell, K. Bruvold, R. Gladden, P. Ilott, and C. Lee, "Relay Support for the Mars Science Laboratory and the Coming Decade of Mars Relay Network Evolution," Proceedings of the 2012 IEEE Aerospace Conference, pp. 1–11, Big Sky, Montana, March 3–10, 2012.

[16] C.D. Edwards, Jr., D.J. Bell, R.E. Gladden, P.A. Ilott, M.D. Johnston, J.L. Maxwell, R. Mendoza, G.W. McSmith, C.L. Potts, B.C. Schratz, J.M. Srinivasan, P. Varghese, "Relay Support for the Mars Science Laboratory Mission", IEEE Aerospace Conference, Big Sky, MT, March 2-9, 2013.

[17] C.D. Edwards, P.R. Barela, R.E. Gladden, C.H. Lee, "Replenishing the Mars Relay Network", IEEE Aerospace Conference, Big Sky, MT, Feb. 2014.

[18] G.J. Kazz, S. Burleigh, K. Cheung, B. Shah, Biren, "Evolution of the Mars Relay Network End-to-End Information System in the Mars Human Era (2030-2040)", IEEE Aerospace Conference, Big Sky, MT. Feb. 2016.

[19] D.D. Wenkert, R.E. Gladden, C.D. Edwards, P. Schmitz, M. Denis, A.J. Winton, "Enabling International Data Relay at Mars," *Space Operations: Contributions from the Global Community*, edited by C. Cruzen et al., Spring, Cham, Switzerland, 2017, pp. 175–205.

[20] "Computer Networks: Crash Course Computer Science #28." Edited by PBS Digital Studios, YouTube, Crash Course, 13 Sept. 2017, www.youtube.com/watch?v=3QhU9jd03a0.

[21] "Computer Networks: Crash Course Computer Science #29." Edited by PBS Digital Studios, YouTube, Crash Course, 20 Sept. 2017, https://www.youtube.com/watch?v=AEaKrq3SpW8.

[22] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, H. Weiss, "Delay-tolerant networking architecture," Internet Engineering Task Force (IETF) Request for comments RFC 4838, April 2007, [Online]. Available: http://www.ietf.org/rfc/rfc4838.txt.

[23] K. Scott, S. Burleigh, "Bundle protocol specification," Internet Engineering Task Force (IETF) Request for Comments RFC 5050, Nov. 2007, [Online]. Available: http://www.ietf.org/rfc/rfc5050.txt.

[24] M. Ramadas, S. Burleigh, S. Farrell, "Licklider Transmission Protocol—Specification," Internet Engineering Task Force (IETF) Request for Comments RFC 5326, Sept. 2008, [Online]. Available: http://www.ietf.org/rfc/rfc5326.txt.

[25] C.E. Edwards, Jr., M. Denis, "Operations Concept for a Solar System Internetwork (SSI)", Interagency Operations Advisory Group (IOAG), Space Internetworking Strategy Group (SISG) report, 15 Oct. 2010.

[26] J. Wyatt, S. Burleigh, R. Jones, L. Torgerson, S. Wissler, "Disruption Tolerant Networking Flight Validation Experiment on NASA's EPOXI Mission," *2009 First International Conference on Advances in Satellite and Space Communications*, Colmar, 2009, pp. 187-196.

[27] A. Schlesinger, B.M. Willman, L. Pitts, S.R. Davidson, W.A. Pohlchuck, "Delay/Disruption Tolerant Networking for the International Space Station (ISS)," *2017 IEEE Aerospace Conference*, Big Sky, MT, 2017, pp. 1-14.